



Release Notes

Release	5.4.0.3
Supported Product	StorNext
Date	April 2017

Contents

What's New in StorNext 5.4.0.3	3
Purpose of this Release	3
Fixed Issues and Enhancements Addressed in StorNext 5.4.0.3	3
StorNext Compatibility	4
Quantum Operating System Upgrade Support Policy	4
StorNext and Linux Interoperability	5
Supported StorNext Upgrade Paths and Upgrade Considerations	5
StorNext Software Upgrade Matrix	5
Considerations for the StorNext File System Directories	5

© 2017 Quantum Corporation. All rights reserved. Your right to copy this manual is limited by copyright law. Making copies or adaptations without prior written authorization of Quantum Corporation is prohibited by law and constitutes a punishable violation of the law. Artico, Be Certain (and the Q brackets design), DLT, DXi, DXi Accent, DXi V1000, DXi V2000, DXi V4000, DXiV-Series, FlexTier, Lattus, the Q logo, the Q Quantum logo, Q-Cloud, Quantum (and the Q brackets design), the Quantum logo, Quantum Be Certain (and the Q brackets design), Quantum Vision, Scalar, StorageCare, StorNext, SuperLoader, Symform, the Symform logo (and design), vmPRO, and Xcellis are either registered trademarks or trademarks of Quantum Corporation and its affiliates in the United States and/or other countries. All other trademarks are the property of their respective owners. Products mentioned herein are for identification purposes only and may be registered trademarks or trademarks of their respective companies. All other brand names or trademarks are the property of their respective owners. Quantum specifications are subject to change.

Journal Size Guidelines	6
Distributed Data Mover (DDM) Guidelines	6
Considerations When Upgrading NFS Server Nodes to StorNext 5	6
Database Schema Update During Upgrades	6
Compatibility Between StorNext and Other Products	7
NAS	7
Infiniband	7
Lattus	7
Partial File Retrieval	8
StorNext Web Services	8
Apple Xsan	8
Supported Browsers	8
General Considerations	8
Checksum Performance Considerations	9
Upgrading Appliances	9
Known Issues	9
StorNext File System Known Issues	10
StorNext Storage Manager Known Issues	12
StorNext GUI Known Issues	15
StorNext Installation, Replication, HA, and Other Known Issues	18
Contacting Quantum	21
Getting More Information or Help	21

What's New in StorNext 5.4.0.3

Purpose of this Release

The StorNext 5.4.0.3 release is a collection of important software fixes that solve several customer-facing issues, including: StorNext file system manager and kernel panics, constant disk rescans and slow cvmkfs performance. Additionally, this release package includes fixes to the StorNext access control (snacl), and Storage Manager recursive retrieve commands.

See the following table for a complete list of the issues addressed by StorNext 5.4.0.3.

Fixed Issues and Enhancements Addressed in StorNext 5.4.0.3

Operating System	Change Request Number	Service Request Number	Description
All	64820	324979	fsm can panic with OPEN_BUSY set in open_deref_nolock when open_refs hits 0.
All	65054	3733044, 317792, 319468, 327212	Constant disk rescans on clients with stripe group marked as down
All	65479	n/a	No signing and authentication for GET operation for AZURE.
All	65595	300308	Reserved inodes aren't always given back when a transaction is aborted
All	65987	n/a	Creates sometimes return EIO
All	66224	314826, 317863, 319270, 321761, 327347, 330014	Deadlock with Open_clients_lock at 5.4.0.1
All	66482	317659	Recursive fsretrieves fail if any file in recursion is on disk. 5.4.0.1 regression.

Operating System	Change Request Number	Service Request Number	Description
All	66502	319178	snac1 cannot properly set READONLY permission on Directory
All	66509	n/a	snac1 -lx should provide the ace_mask
All	66681	319572, 317526, 329425	snprobe: disks in use check should not request information about clients
All	66882	322824	MAC client panic "CVFS ASSERTION FAILED: filecookie == ntohq(replydata->td_file_cookie)
All	67130	n/a	ophang limit panic ... lock inversion between VopOpen and Client_remove
All	67136	n/a	Inside the fsm, a client disconnect can hang on the open_lock ... ophanglimit panic.
All	67224	n/a	Sharing Violation when Global Sharemode is enabled can cause data corruption due to client open reference counts being incorrectly initialized
Linux	64758	3727008, 304461, 304936, 309057, 317498, 317142, 321954, 324079	Xcellis kernel crash due to cvfs_proxy_1stn
Linux	65242	3733306	FSM panic - PANIC: /usr/cvfs/bin/fsm ASSERT failed "free_ip->i_idinode.idi_flags & InodeFlagFree" sn/snfs/fsm/inode.c, line 2302

StorNext Compatibility

For information on StorNext compatibility with operating systems, kernel versions, hardware platforms, drives, libraries, StorNext Appliances, StorNext client interoperability, and other compatibility items, see the [StorNext 5 Compatibility Guide](http://www.quantum.com/sn5docs) available online at <http://www.quantum.com/sn5docs>.

Note: SNAPI and Partial File Retrieval information is provided in separate documents.

Quantum Operating System Upgrade Support Policy

StorNext supports any security or functional bug update that applies to the current StorNext-supported Red Hat update level or SuSE Linux Service Patch. StorNext does not support upgrading to an update level or

service patch beyond the currently supported levels shown in the [StorNext 5 Compatibility Guide](http://www.quantum.com/sn5docs) available online at <http://www.quantum.com/sn5docs>.

StorNext and Linux Interoperability

Newer versions of the Linux `tail` command leverage the `inotify` mechanisms within Linux. The `inotify` mechanisms in Linux are not triggered by file updates coming from other StorNext nodes.

When using the `tail` command on files located in StorNext, Quantum recommends using the following option:

```
---disable-inotify
```

Recommended usage:

```
tail ---disable-inotify -f filename
```

Supported StorNext Upgrade Paths and Upgrade Considerations

StorNext Software Upgrade Matrix

For information on which StorNext versions allow you to upgrade directly to this release, refer to the **StorNext Software Upgrade Matrix** section in the [StorNext 5 Compatibility Guide](http://www.quantum.com/sn5docs) available online at <http://www.quantum.com/sn5docs>.

Considerations for the StorNext File System Directories


On upgrades to StorNext 5, it may be noted that the attributes of many directories in the StorNext file system show much smaller sizes, even zero sizes, where these same directories showed non-zero sizes in StorNext 4.x. This is expected behavior.

Journal Size Guidelines

The absolute minimum Journal Size in StorNext 5 is 4 MB. If a file system is configured with a Journal Size smaller than 4 MB, the Journal Size must be increased prior to upgrading. The recommended Journal Size is 64 MB. New file systems must have a Journal Size of 64 MB or larger.

Distributed Data Mover (DDM) Guidelines

Distributed Data Movers (DDMs) must be upgraded to the same version of StorNext that the Metadata Controller (MDC) is running.

 **WARNING:** Upgrades (such as platform, service pack, etc.) are intended to be done to all systems present in a given deployment. For example, if Xcellis, M660, M440, Pro Foundation, Artico, and G300 are present, they all must be upgraded. One appliance cannot be "left behind".

Considerations When Upgrading NFS Server Nodes to StorNext 5

Due to the fact that the full 64-bit inode numbers are exposed to Linux after Linux clients are upgraded to StorNext 5, special consideration must be made for Linux NFS servers.

To prevent issues with mounted NFS clients, NFS clients must be unmounted prior to upgrading StorNext on the NFS server. If unmounting all NFS clients is not an option during the upgrade, Quantum suggests using the "compat32" mount option on NFS servers.

Database Schema Update During Upgrades

Database schema updates are applied to Storage Manager when upgrading from StorNext 4.3.x and StorNext 4.7.x to StorNext 5. The M660 appliance can achieve approximately one hour for every hundred million entries in the filecomp tables. Smaller appliances and software only configurations may take considerably longer depending on CPU speed and memory availability.

Note: The database schema update conversion time from StorNext 4.7.x to StorNext 5 is significantly faster than that from StorNext 4.3.x to StorNext 5.

StorNext file systems are accessible while the database schema is being updated, but Storage Manager functionality (including stores and retrieves) will be offline.

Do NOT interrupt StorNext services while the database is being updated. Interrupting the database schema update could result in an inconsistent database, and may require assistance from Quantum Support to repair or restore the database.

Use the following commands to determine the number of filecomp entries on the StorNext primary node:

1. List the managed filesystems configured.

```
mysql -e "select Device_key, Path from tmdb.devdb_v;"
```

2. For each *<Device_key>* number listed, display a count of the number of entries in the corresponding filecomp table:

```
mysql -e "select count(*) from tmdb.filecomp<Device_key>;"
```

Note: The query in **Step 2** might require a significant amount of time. Quantum recommends you execute the query before the day of an upgrade.

Compatibility Between StorNext and Other Products

This section describes various interactions between this release and StorNext components and features.

NAS

To view supported StorNext NAS software configurations, see the *StorNext NAS Compatibility Guide* available online at http://www.quantum.com/snnas_cg.

Infiniband

StorNext 5 works with Infiniband SRP (SCSI RDMA Protocol) attached storage for Linux and Windows 2008R2.

Lattus

Refer to the *Lattus Release Notes* for information about compatibility between Lattus and StorNext 5.4.0.3. Object Storage documentation is available online at <http://www.quantum.com/lattusdocs>.

Partial File Retrieval

StorNext Partial File Retrieval (PFR) is a product which enables you to quickly retrieve and utilize segments of large media files, rather than the entire file, based on time-code parameters.

Note: For Quantum Cloud Storage, PFR is not supported for copies with client-side encryption or compression. It is only supported for copies with server-side encryption or without encryption and compression.

For information about compatibility between PFR and StorNext 5.4.0.3, see the *StorNext Partial File Retrieval Compatibility Guide* available online at <http://www.quantum.com/sn5docs>.

StorNext Web Services

StorNext Web Services enables you to run third-party application program interfaces (APIs) with StorNext. To view the latest commands supported by the StorNext Web Services, refer to the [StorNext 5 Web Services Guide](http://www.quantum.com/sn5docs) available online at <http://www.quantum.com/sn5docs>.

Apple Xsan

Xsan is software that enables multiple Mac computers to concurrently access hundreds of terabytes of content on Xserve RAID or Promise RAID storage over high-speed Fibre Channel which allows you to share data faster and consolidate projects. Quantum supplements this solution with StorNext data management software, enabling Apple Xsan customers to use applications running on Windows, Linux, and UNIX with their Xsan and share content across more systems.

For information about compatibility between Apple Xsan and StorNext 5.4.0.3, refer to the [StorNext 5 Compatibility Guide](http://www.quantum.com/sn5docs) available online at <http://www.quantum.com/sn5docs>.

Supported Browsers

For information on browsers supported with the StorNext GUI for this release, refer to the [StorNext 5 Compatibility Guide](http://www.quantum.com/sn5docs) available online at <http://www.quantum.com/sn5docs>.

General Considerations

This section provides information about items to consider for StorNext 5.4.0.3.

Checksum Performance Considerations

Note: Generating MD5 checksums is a CPU intensive operation.

Current StorNext metadata controller and Mover hardware is able to calculate MD5 checksums at around 300 MB/s to 500 MB/s. For newer generation tape technology, the maximum throughput may exceed the rate at which the system can generate checksums. In this case, the MD5 checksum calculation will define the throughput of a single data movement operation. With multiple movement streams, MD5 calculations will be done in parallel across the streams and aggregation of performance will be seen.

Upgrading Appliances

For instructions on upgrading your firmware, refer to the current Release Notes for your particular appliance:

- For Xcellis, see <http://www.quantum.com/xcelliswfdocs>.
- For Artico, see <http://www.quantum.com/articodocs>.
- For M660, M440, M330 Metadata Appliance and Pro Foundation, see <http://www.quantum.com/snmdcdocs>.
- For G300 Gateway Appliance, see <http://www.quantum.com/sngatewaydocs>.

Known Issues

The following sections list known issues in this release of StorNext, as well as associated workarounds, where applicable:

- [StorNext File System Known Issues on the next page](#)
- [StorNext Storage Manager Known Issues on page 12](#)
- [StorNext GUI Known Issues on page 15](#)
- [StorNext Installation, Replication, HA, and Other Known Issues on page 18](#)

Note: If you encounter one or more of the issues listed in this section, please contact Quantum Customer Support and report the issue(s) you encountered. Also inform the support representative whether you were able to successfully work around the issue(s) by using the provided workaround. Doing these things will help Quantum prioritize the order in which known issues are addressed in future StorNext releases.

StorNext File System Known Issues

[Table 1 below](#) lists known issues specific to the StorNext File System process.

Table 1: StorNext File System Known Issues

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	54834	3505208, 3516356	<p>If a file is being copied to the StorNext file system using Windows Explorer and Windows Explorer crashes before it finishes copying all the data, the file may contain data blocks from old, deleted files. This problem occurs because Windows Explorer sets EOF to the size of the file before it writes the data to the file. This leaves a gap of uninitialized data in the file.</p> <p>Note: This problem can also occur with other programs that set EOF beyond the end of data.</p> <p>This problem does not occur if Windows Explorer encounters an error while writing the file; Windows Explorer will delete the partially written file.</p> <p>Workaround:</p> <p>To prevent this problem from occurring on StorNext, you can use the StorNext "client configuration" application's advanced mount option "Restrict Pre-allocation API" on Window systems and the "protect_alloc=yes" mount option on Linux systems. This option will set the unwritten parts of the file to zero. When this option is set, non-root users are unable to use the preallocation ioctl. This option also implies sparse=yes.</p> <p>For more information on this option, see the man page <code>mount_cvfs(8)</code>. The sparse option will introduce some overhead when using Windows Explorer. Before setting the <code>protect_alloc</code> option, see the sparse option in <code>mount_cvfs(8)</code> for a description of how it changes StorNext behavior.</p>
All	64963	n/a	<p>See Conditions for Mounting an Xsan Volume as a StorNext Distributed LAN Client (DLC) below.</p>

Conditions for Mounting an Xsan Volume as a StorNext Distributed LAN Client (DLC)

Beginning with Apple's OS X El Capitan (10.11), support is provided to mount Xsan volumes as StorNext DLC clients. This section provides additional information which should clarify the conditions necessary for this mount to succeed.

When the StorNext services are loaded, the environment is scanned to determine the default mount type - storage area network (SAN) with access to data disks, or DLC with proxy access to data disks through a StorNext gateway node.

If a fibre channel card is detected or internet Small Computer Systems Interface (iSCSI) logical unit numbers (LUNs) are present, the default mode is SAN. If neither of these conditions are true, the default mode is DLC.

In the case that a SAN mount is attempted and not all data LUNs for the file system are present, the mount fails and no error message is logged. If a DLC mount is desired, there are several ways to accomplish this.

First, if the presence of iSCSI LUNs is causing the SAN mount attempt, and these LUNs are not needed, the iSCSI targets can be disconnected. If this is done, the StorNext service must be unloaded and loaded to detect the change.

Example of unload and load:

```
sh-3.2# launchctl unload /System/Library/LaunchDaemons/com.apple.xsan.plist
sh-3.2# launchctl load -w /System/Library/LaunchDaemons/com.apple.xsan.plist
```

If iSCSI LUNs or a fibre channel card are present, the DLC mount is accomplished by indicating this to the driver. There are two ways to achieve this:

- [Set the Xsan Preferences Payload "preferDLC" for the Volume that is to be DLC Mounted below](#)
- [Set the diskproxy=client Mount Options in the /etc/fstab File below](#)

Set the Xsan Preferences Payload "preferDLC" for the Volume that is to be DLC Mounted

Example of setting preferDLC payload and verifying:

```
sh-3.2# defaults read /Library/Preferences/com.apple.xsan
{
  enableSpotlightServer = 1;
}
sh-3.2# defaults write /Library/Preferences/com.apple.xsan preferDLC '(snfs1)'
sh-3.2# defaults read /Library/Preferences/com.apple.xsan
{
  enableSpotlightServer = 1;
  preferDLC = (
    snfs1
  );
}
```

For more information on the preferDLC payload and Xsan configuration profiles, see <https://support.apple.com/en-us/HT205333>.

Set the diskproxy=client Mount Options in the /etc/fstab File

An alternative way to indicate to the driver that a DLC mount is desired is to set the diskproxy=client mount options in the /etc/fstab file.

Example of the diskproxy=client mount options in /etc/fstab file:

```
sh-3.2# cat /etc/fstab
LABEL=snfs1    none    acfs    rw,diskproxy=client
sh-3.2#
```

In either of the above scenarios, a driver unload and load must be performed to enact the change. Once the Volume is mounted, its status as a DLC client can be verified by using the sub-command **who** of the **cvadmin** CLI. First select the file system, then enter **who**. Look for the **C** after the **CLI** in the client entry.

Example of the **cvadmin who** sub-command:

```
Xsanadmin (snfs1) > select snfs1
Xsanadmin (snfs1) > who
Who (File System "snfs1")
```

#	acfs	I.D.	Type	Location	Up Time	License Expires
0>			FSM	smw-rhel64.mdh.quantum.com	4d 23h 33m	N/A
4>	50569B3B7F		CLI S	smw-rhel64.mdh.quantum.com	4d 23h 33m	*10-27-2016 23:59:59
9>			ADM	smw-elcapitan.mdh.quantum.com	0d 0h 0m	N/A
17>			CLI C	smw-elcapitan.mdh.quantum.com	0d 0h 6m	N/A

```

1 active connection out of 65535 licensed connections
0 active connections out of 65535 licensed proxy connections
1 active gateway connection
1 active client-licensed connection

Xsanadmin (snfs1) >
```

StorNext Storage Manager Known Issues

[Table 2 on the next page](#) lists known issues specific to StorNext Storage Manager.

Table 2: StorNext Storage Manager Known Issues

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	43320	1581004	<p>File retrieves from media to disk can be suboptimal for fast tape drives like the Oracle STK T10K drives. This scenario can occur when the retrieve event is initiated on a host that is different from the host running the mover process, which requires the use of synchronous direct I/O.</p> <p>Workaround:</p> <p>To work around this issue and achieve optimal performance for both file stores and retrieves with the T10K drives, increase the default I/O size used by the mover process and make the mover process use asynchronous buffered I/O when the use of synchronous direct I/O is not required, using the following steps:</p> <p>Note: This workaround may also help improve the performance of the faster LTO drives like LTO-6, and LTO-7 by updating the FS_LTO_BLOCK_FACTOR sysparm.</p> <p>Note: Changes to FS_XXX_BLOCK_FACTOR only affects tapes formatted after the change.</p> <ol style="list-style-type: none"> Change the FS_T10K_BLOCK_FACTOR sysparm from 8 to 32 by adding the following entry to <code>/usr/adic/TSM/config/fs_sysparm_override</code>: <pre>FS_T10K_BLOCK_FACTOR=32;</pre> <p>Note: The T10K default I/O block size is 512 KB or 8 * 64 KB. With the block factor changed to 32, the new T10K I/O block size will be 2 MB or 32 * 64 KB. Presently, the FS_T10K_BLOCK_FACTOR sysparm must not be set to a value that exceeds 32.</p> Restart Storage Manager to ensure the change in Step 1 goes into effect: <pre># tsmstop # tsmstart</pre> Verify the FS_T10K_BLOCK_FACTOR sysparm contains the new value: <pre># showsysparm FS_T10K_BLOCK_FACTOR FS_T10K_BLOCK_FACTOR=32</pre> Save the current copies of your <code>/etc/fstab</code> on the MDCs and the DDM clients. Modify <code>/etc/fstab</code> on the MDCs and the DDM clients to use the <code>auto_dma_write_length</code> and <code>auto_dma_read_length</code> mount options as follows: <pre>snfs1 /stornext/snfs1 cvfs rw,auto_dma_write_length=16m,auto_dma_read_length=16m 0</pre>

Operating System	Change Request Number	Service Request Number	Description/Workaround
			<p>0</p> <ol style="list-style-type: none"> 6. Unmount and re-mount your file systems. 7. Use new T10K media to store a copy of the file from the disk. <p>Note: Step 7 is very important; when the new copy is made to the new tapes, the new tapes are labeled with a 2 MB block size, which is used for subsequent writes or reads to and from the media. Tapes on which fsformat was run before the change will use the block factor in use at that time. This change will not impact those tapes.</p>
All	46693	n/a	<p>Executing the command snbackup -s while a full or partial backup is running may result in a message that <code>/usr/adic/TSM/internal/locks/backup.1f</code> is in an invalid format.</p> <p>This is due to the snbackup -s process reading the backup.1f status file while the backup process is updating it.</p> <p>Workaround:</p> <p>Ignore the message; to clear-up the process, re-execute the command snbackup -s (provided that the backup is not writing to the backup.1f status file while snbackup -s is trying to read it again).</p>
All	47833	n/a	<p>When copying files between media using the CLI command fsmedcopy, the file is not re-segmented to match the segment size of the destination media. Rather, the original segments are copied to the target media type and the distribution of segments across destination media will, therefore, be the same as the distribution on the source media.</p> <p>Note: This behavior may cause file data segment distribution to be sub-optimal on the destination media.</p> <p>Workaround:</p> <p>Currently, a workaround does not exist for this Known Issue.</p>
All	65077	n/a	<p>During StorNext Storage Manager start-up, SCSI Reservation conflicts may be reported by the IBM APFO driver if installed, particularly after involuntary fail-over of the StorNext MDC in High Availability configurations.</p> <p>Workaround:</p> <p>These reservation conflicts can generally be ignored and cause no issues with StorNext Storage Manager.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
Linux	65364	n/a	<p>When Storage Manager stores files to LTFS media, the fs_fmover process for TSM generates warning messages in the trace_05 files (found in the /usr/adic/TSM/log/trace directory) indicating that the fs_fmover process is unable to get the extended attributes on new files on the LTFS formatted media. The following is an example of the warning message:</p> <pre> /usr/adic/TSM/logs/trace/trace_05 Oct 25 13:31:49.151125 northern sntsm fs_fmover[22366]: E1201 (8)<02507>:mdt2l1tfs1605: {1}: mdt2l1tfs_file_get_attr: failed to get attr user.QUANTUM_STORNEXT_PATH on file /.StorNext/objects/00/00/01/0000000000000172: -1040 </pre> <p>Workaround:</p> <p>These warning messages are safe to ignore for newly created files on the LTFS media. There may be up to four messages printed per LTFS file, one for each of the StorNext extended attributes on the LTFS media:</p> <ul style="list-style-type: none"> • QUANTUM_STORNEXT_PATH • QUANTUM_STORNEXT_FKEY • QUANTUM_STORNEXT_VERSION • QUANTUM_STORNEXT_META_VERSION

StorNext GUI Known Issues

[Table 3 on the next page](#) lists known issues specific to the StorNext GUI.

Table 3: StorNext GUI Known Issues

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	57856	n/a	<p>The possibility of having different NIC cards installed in the same slots across boots results in having the same Ethernet alias names being used for the network interfaces of different NICs with different speeds (1G/10G).</p> <p>However, the Ethernet alias names depicted in the StorNext Metrics GUI page do not reflect this possible change of the network device representing the alias.</p> <p>Workaround:</p> <p>There is currently no workaround for Change Request Number 57856.</p> <p>Note: StorNext Metrics data is only kept for 30 days; the StorNext Metrics Report auto-corrects any port-discrepancies 30 days after the configuration is changed.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	65328	n/a	<p>When the MDC is acting as a Distributed LAN gateway for a file system and its configuration is modified using the StorNext GUI or the Connect GUI, a spurious LAN server configuration file named <code>/usr/cvfs/config/dpserver.<fsname></code> may be created if it does not already exist. This file will have all network interfaces enabled and overrides the main <code>dpserver</code> file, <code>/usr/cvfs/config/dpserver</code>. In some cases, having all network interfaces enabled is undesirable as it may result in sub-optimal DLC performance and cause DLC traffic to interfere with other MDC network traffic.</p> <p>Workaround:</p> <p>Complete all configurations of the file system prior to setting <code>diskproxy=server</code> in the <code>/etc/fstab</code> file.</p> <p>Or, if the file system configuration requires modification after <code>diskproxy=server</code> has been set in the <code>/etc/fstab</code> file perform the following steps:</p> <ol style="list-style-type: none"> 1. Execute the command <code>sndpscfig -e</code> to verify that <code>/usr/cvfs/config/dpserver</code> exists and contains correct information: <pre># /usr/cvfs/bin/sndpscfig -e</pre> 2. Make <code>/usr/cvfs/config/dpserver.<fsname></code> be a symbolic link to <code>/usr/cvfs/config/dpserver</code>: <pre># /bin/rm -f /usr/cvfs/config/dpserver.<fsname> # ln -s /usr/cvfs/config/dpserver /usr/cvfs/config/dpserver.<fsname></pre> 3. (Optional) Repeat Step 1 through Step 3 for all file systems for which the MDC is acting as a gateway and where <code>/usr/cvfs/config/dpserver</code> is used.

Operating System	Change Request Number	Service Request Number	Description/Workaround
Linux	47954	n/a	<p>The Safari browser becomes unresponsive when you attempt to configure an Email server using the StorNext GUI.</p> <p>Workaround:</p> <p>To workaround this issue, perform the following procedure:</p> <ol style="list-style-type: none"> 1. Shut down the Safari browser window(s). 2. Restart the Safari browser, and then retry operation. 3. Uncheck the Verify SMTP Server Connectivity box, and then retry the operation. 4. Set Authentication to NONE, and then retry operation. 5. Disable the Safari User names and passwords AutoFill under Safari > Preferences > AutoFill, and then retry operation.

StorNext Installation, Replication, HA, and Other Known Issues

[Table 4 on the next page](#) lists known issues specific to StorNext installations, data replication, HA systems, and other areas.

Table 4: StorNext Installation, Replication, HA and Other Known Issues

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	47041	n/a	<p>A database index named <code>classndxatimeme</code> will be automatically added to the <code>tmdb.tier000files%</code> and <code>tmdb.tier001files%</code> tables upon starting TSM for the first time after upgrading from StorNext 5.0.0 or earlier.</p> <hr/> <p>Note: The <code>classndxatimeme</code> index will already be present and not be added again if upgrading from StorNext 5.0.1 or later.</p> <p>Workaround:</p> <p>To minimize TSM downtime after upgrade, the <code>classndxatimeme</code> index can be created prior to performing the upgrade using the <code>index_tierfiles.pl</code> PERL script. This file is available in the StorNext installation directory by opening a support ticket and requesting the file. (Quantum service and service partners can obtain this file from the StorNext Metadata Appliances page on CSWeb.) The script can be run while TSM is running, although it may impact the performance of other operations while the index is being added to the database.</p> <p>To manually add the index, you must have the <code>index_tierfiles.pl</code> script. Then do the following:</p> <p>Log in to the primary server node, and access the command line of the system:</p> <ol style="list-style-type: none"> 1. Log in to the primary server node. 2. Enter the following to source the profile: <pre>./usr/adic/.profile</pre> 3. Change to the directory where <code>install.stornext</code> resides on the installation media. For example: <pre>cd /tmp/stornext/stornext_full/RedHat60AS_26x86_64</pre> 4. Verify that the database is up by running: <pre>mysql_control start</pre>

Operating System	Change Request Number	Service Request Number	Description/Workaround
			5. Execute the PERL script:
			<pre data-bbox="683 401 1458 474">./TSM/index_tierfiles.pl</pre>
			The procedure is complete.
All	57789	n/a	<p>When running StorNext replication from a deduplication-enabled StorNext filesystem to a deduplication-enabled HA StorNext filesystem, corruption is sometimes seen in files on the HA target if HA failovers occur during the replication. Attempts to read the contents of such files report EIO (5), input/output error.</p> <p>Examination of such files using snpolicy's report directive show no TAG or BLK_TAG_PRESENT flags in the file inode.</p> <p>It is not known if this problem is an alternate manifestation of CR 58814, which may also occur in these configurations.</p> <p>Workaround:</p> <p>Avoid HA failovers/reconfiguration while StorNext replication of deduplicated files is active. See also the Release Note item for CR 58814.</p>
All	58814	n/a	<p>When running replication from a deduplicated file system to another deduplicated file system, corruption may occur in large files. Examination of the replicated files using snpolicy's report directive shows that there are holes in the replicated file's extended representation blobmap.</p> <p>These areas read as zeroes rather than the appropriate data.</p> <p>Workaround:</p> <p>Avoid replicating deduplicated files, especially when deduplication is in progress. The problem is more likely to be seen on very large files (tens or hundreds of gigabytes or larger). It is also more likely to be seen when replication is run on a very short schedule (minutes). It is believed to occur mostly or only when deduplication and replication are running concurrently.</p>

Contacting Quantum

Contacts

For information about contacting Quantum, including Quantum office locations, go to:

<http://www.quantum.com/aboutus/contactus/index.aspx>

Getting More Information or Help

For further assistance, or if training is desired, contact the Quantum Customer Support Center:

Region	Support Contact
North America	1-800-284-5101 (toll free) +1-720-249-5700
EMEA	+800-7826-8888 (toll free) +49 6131 324 185
Asia Pacific	+800-7826-8887 (toll free) +603-7953-3010

For worldwide support:

<http://www.quantum.com/serviceandsupport/get-help/index.aspx#contact-support>

Comments

To provide comments or feedback about this document, or about other Quantum technical publications, send e-mail to:

doc-comments@quantum.com